

Epistemic Deontology and the Revelatory View of Responsibility

Timothy Perrine

Metaphilosophy, DOI: 10.1111/meta.12597

Abstract. According to *Universal Epistemic Deontology*, all of our doxastic attitudes are open to deontological evaluations of obligation and permissibility. This view thus implies that we are responsible for all of our doxastic attitudes. But many philosophers have puzzled over whether we could be so responsible. The paper explores whether this puzzle can be resolved, and *Universal Epistemic Deontology* defended, by appealing to a view of responsibility I call the *Revelatory View*. On that view, an agent is responsible for something when it reveals the kind of person the agent is. I explore four ways of developing the *Revelatory View* and argue that none ultimately defend *Universal Epistemic Deontology*.

Key Words: Ethics of Belief; Epistemic Deontology; Epistemic Responsibility; Responsibility for Beliefs; Non-Voluntarist Views of Responsibility;

Is an “ethics of belief” possible? That is, to what degree do agents have obligations that apply to their doxastic attitudes? Many believe we have obligations regarding our inquiries into truth. Such inquiries oftentimes result in doxastic attitudes. But some philosophers maintain stronger views on which there are obligations that apply not just to our actions and inquiries but to our doxastic attitudes directly, independently of the actions that causally contribute to those attitudes.

But if an obligation applies to an agent’s doxastic attitude, then that agent is responsible for that attitude. And a well-known puzzle is whether we are, or could, be responsible for our doxastic attitudes. For our doxastic attitudes are different from our chief paradigm of responsibility, intentional actions. But if we are not responsible for our doxastic attitudes, then a robust ethics of belief might not be possible.

The overall aim of this paper is to evaluate whether one view of responsibility implies that we are responsible for all of our doxastic attitudes and thus makes possible a robust ethics of belief. I call this view the *Basic Revelatory View* because it understands responsibility in terms of revealing the kind of person we are. Nonetheless, I argue that several ways of developing the *Basic Revelatory View* fail to imply that we are responsible for all of our doxastic attitudes. This casts doubt on whether the *Basic Revelatory View* could serve as a foundation for a robust ethics of belief.

The layout is as follows. In section I, I sketch the robust view I’m interested in—*Universal Epistemic Deontology*, the view that any doxastic attitude is open to deontic evaluations, evaluations of obligation or permission. In section II, I indicate that deontic evaluations require responsibility, and clarify the kind of view of responsibility I’m interested in: the *Basic Revelatory View*. However, that view is too underdeveloped to be evaluated. So in sections III–VI, I consider several ways of developing the view, inspired by the work of Robert Adams, David Owens, Mark Heller, and Pamela Hieronymi. I show how each fails to serve as a foundation for *Universal Epistemic Deontology*. In section VII, I distill a fundamental problem for using *Basic Revelatory View* as a foundation for *Universal Epistemic Deontology*. I show how the arguments of the previous sections both illustrate and support this more fundamental problem.

I. Universal Epistemic Deontology

At first approximation, epistemic deontology implies that obligations apply to our doxastic attitudes—our beliefs, disbeliefs, and withholdings.¹ That is, it is appropriate to evaluate doxastic attitudes by saying that an agent *ought* to believe this or *may not* disbelieve that or *shouldn't* withhold on some proposition, etc. By itself, epistemic deontology does not specify the strictness of these obligations. (For instance, perhaps we are never required to believe something but only ever prohibited from believing things.) Nonetheless it implies that our doxastic attitudes are “open” to deontological evaluations, even if different proponents of the position propose different requirements or obligations. While this first approximation leaves open the *scope* of deontic evaluations for doxastic attitudes, here I’m interested in this robust version of epistemic deontology:

Universal Epistemic Deontology: For any agent with doxastic attitudes, any doxastic attitude of that agent is open to deontological evaluations.²

I will briefly mention several reasons to be interested in *Universal Epistemic Deontology*. First, one might just think it is plausible on the face of it! Second, one might think it follows from the essence of doxastic attitudes; doxastic attitude just are the sorts of things that are open to deontological evaluations (see McHugh and Whiting (2014) for relevant discussion). Third, it might be an implication from one’s views about obligations. For instance, some evidentialists maintain that for any doxastic attitude a subject has, one ought to have that attitude if and only if it fits the evidence one has (see, e.g., Feldman (2000), Wood (2008), McCain (2014), Ryan (2015), Stapleford and McCain (2021)). It follows that any doxastic attitude is open to deontological evaluations.

Some embrace weaker views. Consider one way of making precise Williamson’s suggestion that mere belief is “botched knowledge” (2000: 47):

Knowledge Norm for Belief: S ought to believe *p* only if S knows that *p*.

Various authors endorse this principle, or something close to it (see, e.g., Ball (2013: 70-71), Littlejohn (2013), Hindriks (2007: 403), Williamson (2000: 255-6)). This view is weaker than *Universal Epistemic Deontology*, since it lacks implications about withholdings. Nonetheless, those interested in that view may also be interested in *Universal Epistemic Deontology*, since some of the problems of the latter might also apply to the former.

I’ll indicate a further reason that is “meta-normative” in nature. An important issue is which evaluative or normative properties—if any—are basic for epistemology. A basic evaluative property would be one that could be used to analyze or explain other properties. A natural “test” for this is extensional: if the range of things that have one property (e.g. deontic-based) is wider than the range of thing with another (e.g. value-based), then the latter cannot be used to analyze the former.

Rejecting *Universal Epistemic Deontology* might be a first step in an argument that deontic properties are not the most basic for epistemology. Conversely, accepting *Universal Epistemic Deontology* might be a first step in an argument that they are. Of course, those would only be

¹ We could extend the results of this paper to degrees of belief as well. The results of the paper could then be expanded to apply to (e.g.) probabilism.

² The analogous thesis regarding action would be that we are responsible for all our actions. Such a thesis would be widely rejected since we are not responsible for actions caused by (e.g.) external manipulation or phobias. However, I still formulate the position in this way for two reasons. First, it seems to me many epistemologists are interested in the universal or unrestricted form, rightly or wrong. Second, none of my argumentation will turn on cases of (e.g.) external manipulation or phobias. So we could restrict the thesis some and it would not affect my argumentation.

first steps. A lot more argumentation would be needed to develop such arguments. But interest in this meta-normative dispute might, perhaps indirectly, get entangled with *Universal Epistemic Deontology*.

II. From Deontology to Responsibility

If obligations apply to an agent's doxastic attitudes, then the agent is responsible for those attitudes. They are legitimate targets—or “appropriate candidates” as Fischer and Ravizza (1998: 6-7) put it— of various reactions. These reactions may include reactive attitudes such as praise, blame, resentment, appreciation, etc. These reactions may also include other reactions such as rewards or punishment as well. (Compare Fischer and Ravizza (1998: chp. 1), McKenna (2012: chp. 1), Zimmerman (2015).)

This kind of responsibility is what we might call *basic* responsibility. Such responsibility on its own does not imply any “valence.” If an agent is responsible for something in a basic way, this does not tell you whether the agent is praiseworthy, blameworthy, what (if any) reactive attitudes to adopt, etc. To illustrate, I am responsible in a basic way for how my shoes are tied and the small clutter on my desk. Those things can be traced to intentional actions of mind I effectively executed. Nonetheless, I am neither praiseworthy nor blameworthy for those things. Hereafter, when I speak of ‘responsibility’ I have in mind *basic* responsibility.

If *Universal Epistemic Deontology* is true, the agents are responsible for their doxastic attitudes. But responsibility is not a brute, inexplicable feature of reality. So if we are responsible for all of our doxastic attitudes, we can correctly ask *why* we are responsible. A natural way to answer this question is to provide a *view* of responsibility that implies that agents are responsible for all of their doxastic attitudes. Such a view could then be used to defend *Universal Epistemic Deontology*. It would function as a foundation for *Universal Epistemic Deontology*, explaining how it is possible.

Philosophers have proposed various views of responsibility. Some advocate for a *bifurcated* view on the objects of responsibility. According to it, there are certain conditions for responsibility for actions, certain conditions for responsibility for doxastic attitudes, and these conditions are not the same. (Schmitt (1992: chp. 4) and McHugh (2014) are illustrative examples.) By contrast, *unified* views about the objects of responsibility maintain that the conditions for responsibility are the same *regardless* of whether the object of evaluation are actions or doxastic attitudes (or something else). Here I set aside bifurcated views because they raise too many issues for me to discuss.

One traditional, unified family of views of responsibility is voluntarist views. According to these, a necessary condition for responsibility is control or volition of some kind. These views further divided into incompatibilist and compatibilist views of control. According to the former, control is incompatible with determinism; according to the latter, it is consistent with determinism.

Universal Epistemic Deontology is implausible given voluntarist views of responsibility. Given an incompatibilist account of control, *Universal Epistemic Deontology* is false. Alston (1989) famously presents the argument, though it has other adherents. In response to Alston, some have turned to compatibilist accounts of control; see notably work by Matthias Steup (2000, 2008). However, I've argued elsewhere (Perrine (2020)) that given compatibilist accounts of control, *Universal Epistemic Deontology* is false. The upshot is that given a unified voluntarist theory of responsibility, the prospects for *Universal Epistemic Deontology* are dim.

For that reason, I will not examine voluntarist views of responsibility here. Instead, I turn to an entirely different view. It is sometimes called the “real self” view following Wolf (1990). I call it the “Revelatory View.” In general, the view is this:

General Revelatory View: For any ϕ of an agent, that agent is responsible for that ϕ just in case ϕ reveals the kind of person the agent is.

This view can be traced to Hobart (1934) who maintained that agents are responsible for actions that reflect their character, even if agents lack control over their character.

Applying this basic idea to doxastic attitudes gets:

Basic Revelatory View: For any doxastic attitude of an agent, that agent is responsible for that attitude just in case that attitude reveals the kind of person the agent is.

Put so crudely, in such a basic way, *Basic Revelatory View* is in need of development not evaluation. It needs to be developed by giving an account of, first, what makes up “the kind of person” an agent is and, second, what it is for an attitude to “reveal” the kind of person an agent is. Different versions of this view fill in these details differently.

The overall aim of this paper is to see whether *Basic Revelatory View* can be developed to defend *Universal Epistemic Deontology* in a plausible and consistent way. In order for some developed version of *Basic Revelatory View* to defend *Universal Epistemic Deontology* it will have to do two things. First, it will have to get the extension of responsibility right, given *Universal Epistemic Deontology*; that is, it will have to imply that we are responsible for all of our doxastic attitudes. Second, it will have to be a plausible theory of responsibility. Thus, my aim is to see whether *Basic Revelatory View* can be developed to defend *Universal Epistemic Deontology* while doing these two things.

To that end, I will use the work of several philosophers as a stalking horse. I argue that the resulting ways of developing the *Basic Revelatory View* do not defend *Universal Epistemic Deontology* in a consistent or plausible way. Fair warning: my overall aim is not necessarily the same as those authors. My criticisms are not necessarily criticisms of their views or aims. Rather, they are criticisms of potential ways of *using* their views—ways they may or may not be inclined towards. Nonetheless this does not undermine the results. For the *Basic Revelatory View* needs to be filled in. And these authors’ views are potential—and many will think *promising*—ways of doing just that.

III. Robert Adams

In a well-known and frequently cited paper (1985), Robert Adams argues that we are responsible for attitudes, even if we cannot control them. Though his primary aim is to criticize voluntarist views, he does briefly sketch a view of responsibility. He focuses on “states of the mind” that have intentional objects. This would include doxastic attitudes. His sketch of a view appears in the following brief passage:

...we may try to frame a criterion of our accountability for states of mind in terms of the way in which they arise in us. My suggestion is that among states of mind that have intentional objects, the ones for which we are directly responsible are those in which we are responding, consciously or unconsciously, to data that are rich enough to permit a fairly adequate ethical appreciation of the state’s intentional object and of the object’s place in the fabric of personal relationships. Among the states of mind for which we would not be accountable under this criterion are simple feelings of hunger and thirst, insofar as they are primitive responses to physical stimuli... (1985: 26)

Though Adams' view in this passage is nascent, there are still three problems worth mentioning.

First, Adams' view does not contain a plausible development of the idea of "what kind of person" someone is. For Adams' view does not exclude ways of responding to data that fail to reveal what a person is like. Specifically, a person's response to data can be caused and sustained by automatic and subconscious mechanisms that the agent is, and will remain, unaware of. However, responding in such a way does not tell us what kind of person they are. (Beyond the uninteresting sense that they are a person who has mechanisms that responded to data in this way.) A superior account is needed of "what kind of person someone is."

Second, Adams' view of how agents *reveal* what kind of person they are is implausible. Adams focuses on the idea of *responding* to data. But merely responding to data is not sufficient to reveal what kind of person one is. One also needs to be *responsive* to data so that as the data shifts so does the attitude one forms in response. Suppose a person forms an attitude A in response to data D that in some sense "supports" A. (Perhaps the data is someone making a small joke at my expense and the attitude is annoyance with that person.) But suppose that person would still have formed attitude A in response to data D* but data D* does not support forming A. (Suppose this alternative data would have been a joke that was not at anyone's expense.) In the original situation, the person is *responding* to data, but they are not *responsive* to it. A person reveals the kind of person they are by not merely responding to data but by being responsive to it and changing with different data. So Adams' nascent view does not provide a plausible account of revealing either.

These two criticisms indicate that Adams' nascent view is not yet a plausible view of responsibility. But it is doubtful that Adams' view gets the extension right for *Universal Epistemic Deontology*. For it is simply doubtful that *all* of our doxastic attitudes will play an important—or even any—role in our ethical lives and personal relations. Thus, even if Adams' had provided more plausible accounts of what kind of person someone is as well as "revealing," his view would still be inconsistent with *Universal Epistemic Deontology*.

IV. David Owens

David Owens (2000) rejects a voluntarist theory of responsibility. His alternative theory fits the *Basic Revelatory View*. Owens' idea is that responsibility is concerned with one *qua* person (2000: 121, 123). Naturally, he holds that if someone is good *qua* person, then praise is an appropriate response to who they are; whereas, if someone is bad *qua* person, then blame is an appropriate response. But people have many features—they are cooks, lovers, joggers, language users, etc. Which features make up the kind of person they are? Here Owens appeals to people's character traits, their virtues and vices. Further, he is clear, a person's virtues and vices include both ethical and epistemic virtues and vices (2000: 123-4). Such epistemic virtues and vices are character traits that concern how a person "handles" the reasons they have (2000: 126). Owens' examples of epistemic character traits include: being gullible, dogmatic, weak-minded, wise, of good judgment.

Owens maintains that our character traits can be revealed in our thoughts, feelings, and actions (2000: 123). For instance, consider a person with the vice of intemperance (2000: 117). This vice might manifest and reveal itself in an emotional episode of anger. Such an emotional episode might, in turn, cause further emotions or actions, thereby further manifesting the character trait. In this way, on Owens' view, a person might be responsible not just for a particular character trait but also the ways that character trait is revealed or manifested.

Owens' view naturally suggests a way of developing the *Basic Revelatory View*. First, the kind of person one is can be identified with the person's epistemic character traits, their

epistemic virtues and vices. Second, a doxastic attitude reveals the kind of person one is just in case it is a manifestation of that person's epistemic character traits. "Manifesting" is understood in the same way that an episode of anger manifests a vice of intemperance: it is causally generated by the vice and partly composes that vice. Putting these together, we get:

Owens' Condition: For any doxastic attitude of any agent, that agent is responsible for that doxastic attitude if and only if that attitude is a manifestation of that agent's epistemic character traits.

I will raise two problems that show that *Owen's Condition* gets the extension wrong for *Universal Epistemic Deontology*. Both problems involve the requirement that attitudes manifest an agent's epistemic character traits. First, on Owen's approach, an attitude manifests a character trait only if the attitude is part of what it takes to have that character trait. For instance, an episode of anger can reveal a vice of intemperance because a constitutive part of what it takes to have the vice of intemperance is to have dispositions to anger. Such dispositions help distinguish the vice of intemperance from other character traits—like greed, envy, courage, etc. That's why an episode of anger gives us *some* information as to what a person is like or not like, what their character traits are.

However, it is not part of some specific epistemic character trait—some epistemic virtue or vice—that one is disposed to form particular kinds of doxastic attitudes. For instance, one can form a belief by being either being intellectually careful or intellectually sloppy or somewhere inbetween. Thus, forming specific doxastic attitudes is not what distinguishes between different epistemic character traits. Consequently, forming a particular attitude does *not* manifest an epistemic character because, in general, such an attitude is not part of having that character trait. (This is true *even if* that attitude is the causal result of an epistemic character trait.) Thus, in general, doxastic attitudes are not manifestations of an agent's epistemic character traits. For this reason, many doxastic attitudes do not meet *Owen's Condition*.

A second problem is that an attitude manifests a character trait only if it is caused by having that character trait. But epistemic character traits do not play a causal role in generating many doxastic attitudes. Most of our beliefs simply aren't caused by wisdom, intellectual sloppiness, dogmatism, etc. Thus, most of our doxastic attitudes are not manifestations of our epistemic character traits. Thus, most of our attitudes will not meet *Owens' Condition*. Thus, *Owens' Condition* cannot get the extension right and thus serve as a foundation for *Universal Epistemic Deontology*.³

V. Mark Heller

Heller's (2000) view of responsibility for doxastic attitudes explicitly takes inspiration from Hobart. And Heller maintains that a person is responsible for his doxastic attitudes when they reveal the kind of person they are. Thus, I will examine his view despite some potential classificatory difficulties.⁴

³ Recall: Owens understands epistemic virtues and vices like character traits. Other philosophers have defended an approach to epistemic virtues and vices were they are closer to subconscious reliable belief forming mechanism. But that approach isn't Owens' and will likely fall prey to the problems discussed with Adams' view.

⁴ The difficult is that Heller endorses something close to *Basic Revelatory View*—which I presented as a rival to voluntarist views—but he also maintains that responsibility requires control—which I presented as part of voluntarist views. The resolution of the difficulty is that Heller believes that the ordinary view of responsibility is both voluntarist *and* incoherent (2000: 132). So he opts to retain a voluntarist requirement for responsibility while offering a revisionary view of control that is similar to the requirements of *Basic Revelatory View*. Regardless of

A. *Heller's Theory, in Brief*

Heller's view has the following condition:

Heller's Condition: A agent is responsible for a belief only if the belief reveals the agent's epistemic nature.⁵

To further develop *Heller's Condition*, I will exposit Heller's views of what an agent's "epistemic nature" is and when a belief "reveals" that epistemic nature.

In discussing the question of epistemic nature, I will use a three-fold distinction. The distinction is normally marked with regards to beliefs (see, e.g., Audi (1994)), but it applies to other attitudes as well. First, we have occurrent attitudes—attitudes we have that we are currently aware of. Second, we have non-occurrent attitudes—attitudes we have that we are not currently aware of. Finally, we have dispositional attitudes. Despite the label, these are not attitudes. Rather, they are dispositions to form attitudes. When these dispositions manifest, the result is an attitude—either occurrent or non-occurrent.

According to Heller, an agent's epistemic nature is composed of their "desires" (2000: 135ff.). Specifically, these desires are higher-order desires to use certain dispositions, mechanisms, or patterns of reasoning to form or retain beliefs (cf. (2000: 135)). The desires can be occurrent or non-occurrent. Epistemic natures can also include dispositional desires that would cause agents to have occurrent desires upon "appropriate prompting" (2000: 135). I'll refer to all of these as 'higher-order desires' for ease of reference.

According to Heller, a belief reveals an epistemic nature only if it is *caused* by that epistemic nature (2000: 133). A belief reveals an epistemic nature "just in case S's desires to have certain epistemic dispositions play the appropriate causal role in producing the belief in question" (135-6). "Appropriate causal role" requires more than merely excluding bizarre deviant chains philosophers are fond of creating. A higher-order desire plays an appropriate causal role when the higher-order desire to use a certain disposition (e.g. *modus ponens*) causes one to actually use *that* disposition to create a belief (cf. 2000: 136).

Thus, we can develop *Heller's Condition* as follows:

Heller's Condition Developed: A agent is responsible for a belief only if that agent's higher-order desires to have certain belief-forming dispositions cause the agent to have that belief (in the appropriate way).

B. *Critical Discussion*

I'll raise three issues. First, what Heller identifies as an agent's epistemic nature has implausible results. Second, agents frequently lack the relevant higher-order desires. Finally, even if the agents have the relevant higher-order desires, it is unlikely that the causal condition in *Heller's Condition Developed* is met.

First, Heller identifies an agent's epistemic nature with her higher-order desires to utilize certain belief-forming dispositions. These include occurrent and non-occurrent desires, but also dispositional desires. The inclusion of dispositional desires is problematic for the following reason. *Which* dispositional desires an agent has is partly determined by how reflective the agent is. For a dispositional desire is a disposition that, upon certain prompting, would manifest in a

how best to classify his view, I'll focus on whether it is plausible and consistent with *Universal Epistemic Deontology*.

⁵ Heller assumes responsibility requires control and says that a doxastic attitude is under an agent's control if and only if it reveals their epistemic nature (2000: 135, 136). I interpret Heller's theory about formed and retained doxastic attitudes as this is a natural interpretation and any weaker one wouldn't be consistent with *Universal Epistemic Deontology*.

desire. But whether certain prompting would manifest in a desire can be influenced by how reflective an agent is. However, once we realize this, there may be *too many options* to choose to be *the* agent's epistemic nature. (Is it the ones that result given no reflection, a little, a lot, or maybe "sustained reflection," "reasonable reflection," "socially expected reflection"?) Call this the 'Selection Problem.'

Heller is aware of the Selection Problem and has formulated a response. Which dispositional desires are part of an agent's epistemic nature at a given time are those the agent has at that time given how reflective the agent is at that time. Thus, if *right now* an agent is not very reflective, and as a result has a dispositional desire to condone hasty generalization, then *right now* that dispositional desire is part of their epistemic nature. At a later time, when they are more reflective, they might not have that dispositional desire at that time.

Heller's response has a problematic result for a theory of responsibility. How reflective an agent is at a time can be influenced by an enormous range of factors, such as how interested, tired, excited, emotional, busy, etc. the agent is. Thus, those range of factors are relevant to what epistemic nature an agent has and, by extension, what doxastic attitudes the agent is responsible for. And this is counterintuitive. We have a large number of doxastic attitudes, though most are non-occurrent most of the time. It is perplexing to think that which attitudes I am responsible for could change from moment to moment depending upon how my higher-order desires might change. To put it crudely, whether or not I'm responsible for some beliefs at a time shouldn't be determined by whether (e.g.) I'm hung over or drunk, just went for a run, am sleepy, had too much coffee, am in a foul mood, etc. So given Heller's theory of epistemic nature and his response to the Selection Problem, *Heller's Condition Developed* has implausible results.⁶

Additionally, given *Heller's Condition Developed*, agents are not responsible for all of their doxastic attitudes. Thus, it is not consistent with *Universal Epistemic Deontology*. I'll provide two reasons for thinking that, given *Heller's Condition Developed*, agents are not responsible for all of their attitudes: one concerning a lack of relevant desires, one concerning lack of right causal connection.

According to *Heller's Condition Developed*, an agent is responsible for a belief only if the agent's higher-order desires to have certain belief-forming dispositions causes the agent to have that belief (in the appropriate way). And that occurs only if the disposition that creates the belief is the same as the object of the higher-order desire. Thus, an agent is responsible for a belief only if the agent has a desire to form beliefs in the actual way that the belief is formed. There needs to be a "match" between the desire and actual belief formation.

Presumably, some of the time our higher-order desires match the way we form our beliefs. Heller's chief example of a belief-forming disposition is using the logical rule *modus ponens*. Clearly some have a higher-order desire to use this disposition and do use it from time to time. (Though I do wonder how many people have these specific desires who haven't taught a critical thinking course.) But reasoning in accordance with explicitly formulated logical rules is not representative of all the reasoning we do.

Let's distinguish between belief forming dispositions that are inferential—taking other beliefs as inputs—and belief forming dispositions that are non-inferential. The literature on biases and heuristics describes various kinds of inferential belief-forming dispositions. (Well-known works include Kahneman, Slovic, Tversky (1982) and Kahneman (2011).) However,

⁶ Heller could exclude dispositional desire as part of epistemic nature, restricting an epistemic nature to only occurrent or non-occurrent desires. This would push his position closer to Hieronymi's, which I'll discuss in the next section.

many people are unfamiliar with this literature. As a result, there is not a “match” between their higher-order desires and the way they form beliefs. Similar points will apply to non-inferential belief-forming dispositions. For instance, people have various dispositions to form beliefs as a result of various sensory information. These dispositions are studied in various ways by psychologists, cognitive scientists, and others. But most people are unfamiliar with such work. Consequently, there will not be a match between their higher-order desires and the way they actually form beliefs.⁷

I’ve argued briefly that people lack the relevant kinds of higher-order desire to always “match” the ways they form beliefs. But even if there is a match between the desires, it is doubtful that they play the right *causal* role. For many of our belief-forming dispositions do not require a higher-order desire to use the disposition as a requirement for that disposition to manifest. Indeed, such desires are probably epiphenomenal.

Heller disagrees. He thinks it’s very easy for our higher-order desires to cause our beliefs. (In fact, he even worries it might be too easy, cf. (2000: 138-140).) He gives the following example (2000: 136), writing

I reason by *modus ponens* precisely because I recognize its value—precisely because I endorse that reasoning process. I sometimes reason by hasty generalization because, at the time that I so reason, I am not being reflective enough, and I fail to see that I am reasoning by hasty generalization—I endorse the reasoning pattern that I am using as some acceptable pattern rather than the unacceptable pattern that it is.

However, this example is unconvincing. First, this case involves an *occurrent* desire to use an *inferential* disposition that encodes a *well-understood rule of logic*. Many cases of belief formation fail to have each, if not all, of these features. Further, this case is unlike cases of retained belief as opposed to newly formed beliefs. So this case isn’t representative of all of the various and sundry ways we form and retain beliefs. Second, I’m skeptical that in this case the *occurrent* desire plays any important causal role. More likely, a *modus ponens* disposition was placed in a circumstance where it manifested and produced a certain belief. Perhaps that person has an *occurrent* desire to reason using *modus ponens* in that case. But, again, such a desire is likely epiphenomenal for the disposition to manifest. (If one remains doubtful, consider another case. Someone stands in front of a tree and announces “I’m forming the belief that this is a tree because I recognize the value of perception!” Would anyone really think they have correctly identified the causal origin of this belief?) Thus, many—nonetheless all—of our doxastic attitudes do not meet *Heller’s Condition Developed*. Thus, *Heller’s Condition Developed* gets the extension wrong for *Universal Epistemic Deontology*.

VI. Pamela Hieronymi

Finally, I will discuss Pamela Hieronymi’s sophisticated view, focusing primarily on her (2008) paper aptly entitled “Responsibility for Believing,” though I will utilize her (2014) paper when relevant. In discussing her paper, I will focus on how it can develop the *Basic Revelatory View*. While her paper contains one way of developing that view, it also contains a second line of

⁷ Higher-order desires may include dispositional higher-order desires. Their inclusion does not help. For—setting aside the Selection Problem—it is unlikely that people even have dispositional higher-order desires that match their actual belief forming dispositions. To illustrate, people don’t have dispositional higher-order desires about probabilistic fallacies because people don’t even understand the fallacies to begin with. And generally one does have a disposition about X if one doesn’t understand X.

thought involving “answerability” as sufficient for responsibility. Crudely put, one is answerable for something just in case it would be appropriate for one to *answer* for it, that is, to defend or justify that thing. I won’t discuss this line of thought for two reasons. First, I already discuss answerability views of responsibility in connection to responsibility for doxastic attitudes in a companion paper (Perrine (2022)). Second, I am not sure that Hieronymi’s own way of developing an answerability view is the most plausible for reasons I’ll put in a footnote.⁸

A. Hieronymi’s View

Hieronymi writes things that are reminiscent of the *Basic Revelatory Thesis*:

To morally assess someone on account of some action or attitude...is not just to make a statement about the action or attitude; it is to make a statement about the quality of the will or mind or moral personality of [which the] action or attitude is a product or part (2008: 362).

To morally assess someone on account of some state of affairs... one needs to think that the state of affairs in some ways reveals something of the person, of the quality of her will. (2008: 362)

To use Hieronymi’s work to develop the *Basic Revelatory Thesis* we need to ask two questions. First, using her terminology of ‘quality of mind’ instead of “kind of person,” what account of “quality of mind” does she provide? Second, when does something “reveal” a person’s quality of mind?

First, Hieronymi describes one’s quality of mind as one’s “take on the world,” what is “important or worthwhile or valuable in [the world]” (2008: 361). As she sees it, one’s will is composed of psychological states or attitudes such as beliefs, intentions, or desires. These will frequently be states ascribing normatively salient properties—being worthwhile, important disgusting, etc.—to things like actions or attitudes.

Second, though Hieronymi frequently invokes the idea of revealing one’s quality of mind, she doesn’t provide a detailed account in her (2008) paper. However, an account can be distilled from various passages, especially those in a later paper (2014: 12ff.). I will explain the account using a three-fold distinction. First, an agent’s *reasons for acting* are those things that causally explain why the agent performed that action. On a standard view, reasons for acting will be psychological states such as beliefs and desires. Second, a *reason in favor or against* the agent’s action is anything that is a consideration in favor or against that action. On a standard view, these are not necessarily psychological states. Finally, an *agent’s reason* is what the agent takes to be in favor or against an action. Agent’s reasons are what agents might use to defend or justify their actions.

Clearly the first two categories are distinct. Additionally, the third is distinct from each of them. First, what explains why an agent performed an action need not be what the agent takes to be in favor of that action. This is illustrated in cases of weakness of will. But more mundane cases illustrate the same point. Each morning I get up and stumble to the fridge for water. There are considerations that I take to favor this action. But they are not part of the best causal

⁸ Hieronymi (2008) understands answerability in terms of settling questions. So one is responsible for believing *p* because one has settled the question of whether *p* (cf. (2008: 359-63)). But Hieronymi offers a metaphysically thin understanding of settling the question of whether *p* so that it is equivalent to believing *p* (2008: 360, fn. 5, 360 fn. 8; 2014: 14 fn. 24). This comes perilously close to saying that one is responsible for believing *p* just because one believes *p*! This strikes me as explanatorily defective. Indeed, it seems closer to a *primitivist* view of responsibility for doxastic attitudes. Since I am interested in a fuller explanation for why we might be responsible, I will not focus at any more length on this metaphysically thin development of the answerability view.

explanation of my behavior—a well-engrained habit is frequently a better explanation. Second, what are actual considerations in favor of (or against) an action need not be what the agent takes to be in favor of (or against) an action. Human beings are fallible.

I introduced this distinction for actions. But it will need to be extended to attitudes—including doxastic attitudes—to be relevant for this discussion. Hereafter, I will assume that it can coherently be extended and use the term in this more expansive way.

Returning to the question of “revealing” a quality of mind, something—an action, attitude, etc.—*reveals* a person’s quality of mind when that person has an agent’s reason for that thing (cf. Hieronymi (2014: 12ff)). To illustrate, suppose a person purchases free-trade, organic coffee. His agent’s reason for this might be that it is important to support free-trade organic farming practices. Another person might purchase a brand that is not free-trade or organic. Her reason might be that the resulting coffee is better tasting. These agent’s reasons reveal what the respective individuals value and find worthwhile.

Having an agent’s reason thus “links” actions, attitudes, etc. to those things that the agent finds important, valuable, worthwhile, etc. Further, these “links” are just further attitudes the agent has that are part of the agent’s quality of mind. For instance, in the case of the free-trade, organic coffee buyer, the linking attitude might be a *belief* that purchasing free-trade, organic coffee contributes to an important cause.

On this view, then, what an agent “is like” is identified with their quality of mind which is in turn identified with various attitudes such as beliefs, desires, and perhaps intentions. Additionally, having an agent’s reason is sufficient to “reveal” that quality of mind. This is because one has an agent’s reason only if there is some attitude that is part of the quality of the mind that “links” that thing to what the agent takes to be important, worthwhile, significant, etc. Applying this general view to doxastic attitudes gives us the following view:

Hieronymi’s Condition: If an agent has an agent’s reason for a doxastic attitude, then the agent is responsible for that attitude.

Hieronymi’s Condition (partly) develops the *Basic Revelatory View* because it does not state necessary and sufficient conditions, merely sufficient ones (cf. Hieronymi (2014: 28)). Nonetheless, this will be adequate for my purposes.

Finally, this view improves on previous authors. First, unlike Adams, it explicitly identifies what agents are like with their attitude thereby excluding merely subconscious processes. Second, unlike Owen’s view, it does not require that an attitude reveals what an agent is like only if the agent has a character trait that the attitude is a part of it. (Indeed, Hieronymi is explicit that what an agent is like is not to be identified with character traits (2008: 361 fn. 10).) Finally, unlike Heller’s view, it does not require that the agent’s reason for something needs to “match” what causes the agent to have or retain that attitudes. Nonetheless, despite these improvements, some difficulties remain, or so I’ll now argue.

B. Degrees of Revealing an Agent’s Quality of Mind

First, *Hieronymi’s Condition* is implausible as a sufficient condition for responsibility because it makes too much out of too little. To see this, notice that on some issues an agent’s quality of mind might be quite expansive. An agent might have spent many years ruminating over various kinds of actions and attitudes resulting in many sophisticated opinions. However, an agent might not. She might have some sort of agent’s reason for an attitude or action. But it might be inchoate or undeveloped. And, of course, she might not have any agent’s reason whatsoever on some issues.

However, *Hieronymi's Condition* implies that a transition from no agent's reason to having an agent's reason is *sufficient* for one to be responsible. But given how minimal, weak, inchoate, or underdeveloped an agent's reason might be, this is implausible. An agent's reason could be as simple as "it seemed right" or "I guess that's what they told me." And thus acquiring these agent's reasons, when before one had none, would be all that is required for one to be now be responsible for something. I find this deeply surprising and perplexing.⁹

C. *Scope and Extent of Agent's Reasons*

There is no doubt that for many of our actions and attitudes we have agent's reasons. This is especially clear in cases where agents engage in a deliberative process of weighing considerations in favor or against various attitudes. After such deliberating, agents will frequently have agent's reasons for their attitudes.

However, many of our attitudes—including doxastic attitudes—are not formed in a deliberative process where we weigh considerations. Further, when we form attitudes in a non-deliberative way, it is unlikely that we have agent's reasons or the relevant quality of mind required by agent's reasons. Indeed, many of the attitudes that make up a quality of mind that could be used to link things to agent's reasons are sophisticated attitudes (e.g. attitudes about what are valuable, significant, worthwhile, disgusting, etc.). And the default is that we do *not* form such attitudes unless an occasion calls for it. Deliberating is such an occasion, but most occasions are not.

To make this criticism more specific let's consider two types of cases. First, consider some doxastic attitudes I form immediately and unreflectively. For instance, as I sit here on the plane, I believe that my right knee hurts, I am slightly tired, the plane is full and my coffee is getting cold; I withhold on whether the flight is mostly over, whether they about to collect trash; and I disbelieve that the person next to me needs to use the restroom, and the ambient light is sufficient for reading. These attitudes are formed without some *secondary* attitude linking them to things I find important, worthwhile, disgusting, etc. Because they are formed without such a secondary attitude, it is highly improbable that I have any agent's reasons for them.

Second even if a belief is formed in a reflective way with an agent's reason, there is no guarantee that one *retains* that agent's reason. Here's a well-known example from Alvin Goldman updated for my purposes. Suppose an agent comes to believe that a serving of broccoli is a good source of calcium. When the agent formed their belief, they did so with an agent's reason. (Perhaps they formed the believe by consulting a trustworthy news sources which they explicitly recognized as such.) Nonetheless, in the interim, the agent might have forgotten the origin of their belief and with it their agent's reason. But the belief persists. Doxastic attitudes like this lack agent's reason. Though I don't know how to estimate this, I imagine many doxastic attitudes are like this.

In considering these cases it is important to keep in mind the distinctions introduced in my discussion of Heller's view. My criticism here is not that agents in these cases lacks *occurrent* mental states that would link their attitudes with their agent's reasons. That is, my criticism is not that the relevant quality of mind is not occurrent when they form their attitudes. Rather, my criticism is that they lack the relevant quality of mind to begin with. To be sure, agents might have various dispositions to have attitudes that could be part of their quality of mind. But

⁹ This criticism does assume that agents sometimes do not have agent's reasons, but that assumption is quite plausible. This criticism is consistent with Hieronymi's insistence that one can be rightly asked for the reason for something even if one did that thing "for no reason" (2008: 360; 2014: 13-14). That's because, when she writes that, she has in mind an agent's *reason for acting* not an *agent's reason* (as I use these terms).

dispositions to have attitudes are not themselves attitudes and are not part of an agent's "take" on the world.

I've identified two types of cases where it is plausible that agents will frequently lack agent's reasons for their doxastic attitudes. Thus, *Hieronymi's Condition* will not imply that agents are responsible for their doxastic attitudes in many of those cases. To be sure this does not mean that *Hieronymi's Condition* is inconsistent with *Universal Epistemic Deontology*. But it does mean that proponents of the latter position cannot use the former to defend *Universal Epistemic Deontology*.

VII. A Fundamental Problem

So far I have examined four different accounts that develop the *Basic Revelatory View*. I've argued that none of them can serve as a foundation for *Universal Epistemic Deontology*. They either get the extension wrong—implying that we are only responsible for some of our doxastic attitudes—or are implausible theories of responsibility or both. But my choice of these views was not arbitrary. I partly choose them because of their prominence. But I also partly choose them because they help illustrate a fundamental problem with using the *Basic Revelatory View* as a foundation for *Universal Epistemic Deontology*. In this section, I will describe the fundamental problem.

The *Basic Revelatory View* is the result of applying the *General Revelatory View* to doxastic attitudes. That position was:

General Revelatory View: For any ϕ of an agent, that agent is responsible for ϕ just in case ϕ reveals the kind of person the agent is.

Like the *Basic Revelatory View*, the *General Revelatory View* needs to be developed in order to be plausible. But there are two natural and plausible ways to develop it and, by extension, instances of it, like *Basic Revelatory View*.

First, an account needs to be provided of "the kind of person" the agent is by identifying some feature of the agent. However, whatever feature it identifies, it needs to be sensible that we hold an agent responsible for that kind of thing. For instance, it does not make much sense to hold me responsible for various physiological features (though it may be sensible to hold me responsible for actions that led to those physiological features). So it would be implausible to develop the *General Revelatory View* by claiming that sometimes when I am responsible for something it reveals some physiological feature.

However, it is frequently sensible to hold agents responsible for their attitudes (or something close to them, such as dispositions or manifestations of them). In fact, we already *do* hold people responsible for these kinds of things! Further, differences in the content of attitudes may make sensible different ways of holding agents responsible. Thus, it is natural and plausible to provide an account of "the kind of person" an agent is that identifies it with the agent's attitudes (or something close to them).

Second, an account needs to be provided that explains why some ϕ is linked to some particular feature of the agent as opposed to some other feature. After all, agents have many features. Why does ϕ reveal some feature of the agent but not some other? Here it is natural and plausible to appeal to some internal psychological connection—either attitudinal or causal (or both). That is, it is natural and plausible to think that ϕ reveals some feature of the agent because either ϕ is causally sensitive to ϕ or because the agent herself would link, via some attitude, ϕ and the feature that ϕ reveals.

Bringing these two points together, it is natural and plausible to develop the *General Revelatory View* so that "what the agent is like" is identified with attitudes (or something close to

them) and some ϕ reveals such attitudes when there is an internal psychological connection between ϕ and those attitudes. However, these are not trivial psychological requirements. And, depending upon what the ϕ is under discussion, there is no guarantee that they will be met.

It is at this point that the fundamental problem appears. *Basic Revelatory View* is the result of applying the *General Revelatory View* to doxastic attitudes. Thus, on the one hand, developing the *Basic Revelatory View* so that it has these non-trivial psychological requirements makes it a more plausible theory of responsibility. But, on the other hand, the more non-trivial psychological required for responsibility, the fewer doxastic attitudes agents will be responsible for. So the fundamental problem is that it seems as if one can develop the *Basic Revelatory View* so that it is more plausible *or* consistent with *Universal Epistemic Deontology* but *not* both.

I have briefly presented this fundamental problem. But my presentation is brief and the problem unadorned. For these reasons, we might not be confident that the problem is legitimate. It may be that when we turn to more specific and concrete ways of developing the *Basic Revelatory View* the problem will dissipate. Crude and vague “problems” oftentimes dissipate when confronted with specific and detailed solutions.

But I *have* looked at several more specific ways of developing the *Basic Revelatory View*. And none of them avoid the fundamental problem. Each either fails to be plausible as a theory of responsibility or fails to have the right extension or both. For instance, of the four theories, I regard Owen’s and Hieronymi’s as, comparatively, the more plausible. But they are more plausible exactly because they have more rigorous psychological requirements; and it is the more rigorous psychological requirements that keep them from getting the extension right. Conversely, Adams’ and Heller’s views may imply that we are responsible for more doxastic attitudes than Owen’s and Hieronymi’s, but they are comparatively less plausible.

In this section I have briefly articulated a fundamental problem with *Basic Revelatory View* being a foundation for *Universal Epistemic Deontology*. And in sections III-VI, I looked at particular ways of developing the *Basic Revelatory View*. And I argued each failed as a foundation for *Universal Epistemic Deontology*. The relationship between the argument in this section and those is two-fold. First the arguments of each of those sections, individually, helps *illustrate* the fundamental problem in this sections. Second, the arguments of those sections, together, *support* the claim that the fundamental problem in this section is not illusory, that it is a genuine problem.

To be sure, my argument that *Basic Revelatory View* cannot be a foundation for *Universal Epistemic Deontology* is not deductive. Rather, it is more like the kind of argumentation we find in the sciences. In the sciences, theorists start with some crude and unsophisticated core claim—say, light is a wave, the sun revolves around the earth, hereditary features are caused by a common cause, etc. These core claims are then evaluated by constructing more specific theories of them. After all, if the core claim is true, then there is *some* way of constructing a more specific theory that is also true. But if several ways of constructing a more specific theory turn out to be false, this is evidence against the core claim.

Similarly, if the *Basic Revelatory View* could serve as a foundation for *Universal Epistemic Deontology*, then there is some more specific version or development of that view that is plausible and consistent with *Universal Epistemic Deontology*. I have argued that several more developed accounts are not. This casts doubt not just on each of those developed accounts being a foundation for *Universal Epistemic Deontology*. It also casts doubt on the more general claim that the *Basic Revelatory View* could serve as a foundation for *Universal Epistemic Deontology*.

VIII. Concluding Remarks

An ethics of belief maintains that our doxastic attitudes are open to deontic evaluations directly, not in virtue of our actions. One robust ethic of belief is *Universal Epistemic Deontology* which maintains that *all* of our doxastic attitudes are open to deontic evaluations. But this position would thereby imply that we are responsible for all of our doxastic attitudes, and it is not clear how we are responsible for doxastic attitudes. In this paper, I've explored whether one view—the *Basic Revelatory View*—could serve as a foundation for *Universal Epistemic Deontology*. I developed this basic view in four different ways and argued that none provide a foundation for *Universal Epistemic Deontology*. I then argued that there is a fundamental problem with a revelatory view being a foundation for *Universal Epistemic Deontology*. Proponents of *Universal Epistemic Deontology* should likely look elsewhere to shore up the foundations of their views.¹⁰

Bibliography:

- Adams, Robert. (1985). "Involuntary Sins." *The Philosophical Review*. 94.1: 3-31.
- Alston, William. (1989). *Epistemic Justification*. Ithaca: Cornell University Press.
- Ball, Brian (2013). "Knowledge is Normal Belief." *Analysis*. 73.1: 69-76.
- Feldman, Richard. (2000). "The Ethics of Belief." *Philosophy and Phenomenological Research*. 60.3: 667-695.
- Fischer, John Martin and Mark Ravizza. (1998). *Responsibility and Control*. Cambridge University Press.
- Heller, Mark. (2000). "Hobartian Voluntarism: Grounding a Deontological Conception of Epistemic Justification." *Pacific Philosophical Quarterly*. 81: 130-141.
- Hieronymi, Pamela. (2008). "Responsibility for Believing." *Synthese*. 161.3: 357-373.
- Hieronymi, Pamela. (2014). "Reflection and Responsibility." *Philosophy and Public Affairs*. 42.1: 3-41.
- Hindriks, Frank (2007). "The Status of the Knowledge Account of Assertion." *Linguistics and Philosophy*. 30: 393-406.
- Hobart, R. E. (1934). "Free Will as Involving Determination and Inconceivable Without It." *Mind*. 43.169: 1-27.
- Kahneman, Daniel (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus, and Giroux.

¹⁰ For helpful comments on a previous draft, I thank Michael Longenecker, Timothy O'Connor, and an anonymous reviewer. Thanks, as well, to a reviewer of my (2020) for encouraging me to examine these views of responsibility more closely in connection with these issues.

- Kahneman, Daniel, Paul Slovic, and Amos Tversky. (1982). *Judgment Under Uncertainty*. Cambridge: Cambridge University Press.
- Littlejohn, Clayton (2013). "The Russellian Retreat." *Proceedings of the Aristotelian Society*. 83:293-320.
- McCain, Kevin. (2014). *Evidentialism and Epistemic Justification*. Routledge Press.
- McHugh, Conor and Daniel Whiting. (2014). "The Normativity of Belief." *Analysis Reviews*. 74.4: 698-713.
- McKenna, Michael. (2012). *Conversation and Responsibility*. Oxford University Press.
- Owens, David. (2000). *Reason Without Freedom*. London: Routledge.
- Perrine, Timothy. (2020). "Strong Internalism, Doxastic Involuntarism, and the Costs of Compatibilism." *Synthese*. 197: 3171-3191.
- Perrine, Timothy. (2022). "Responsibility, Epistemic Deontology, and the Rational Relations View." Unpublished Manuscript.
- Ryan, Sharon. (2003). "Doxastic Compatibilism and the Ethics of Belief." *Philosophical Studies*. 114.1/2: 47-79.
- Ryan, Sharon. (2015). "In Defense of Moral Evidentialism." *Logos & Episteme*. 6.4: 405-427.
- Schmitt, Frederick. (1992). *Knowledge and Belief*. Routledge Press.
- Stapleford, Scott and Kevin McCain (2021). "Bound By Evidence." In *Epistemic Duties: New Arguments, New Angles*, eds. Kevin McCain and Scott Stapleford. Routledge.
- Steup, Matthias. (2000). "Doxastic Voluntarism and Epistemic Deontology." *Acta Analytica*. 15.24: 25-56.
- Steup, Matthias. (2008). "Doxastic Freedom." *Synthese*. 161.3: 375-392.
- Williamson, Timothy. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Wolf, Susan. (1990). *Freedom Withing Reason*. New York: Oxford University Press.
- Wood, Allen. (2008). "The Duty to Believe According to the Evidence." *International Journal for Philosophy of Religion*. 63.1: 7-24.
- Zimmerman, Michael. (2015). "Varieties of Responsibility." In *The Nature of Moral Responsibility*, eds. Randolph Clarke, Michael McKenna, and Angela Smith. Oxford: Oxford University Press.